Fibonacci Heaps Revisited

Haim Kaplan* Robert E. Tarjan† Uri Zwick‡

arXiv:1407.5750v1 [cs.DS] 22 Jul 2014

Abstract

The Fibonacci heap is a classic data structure that supports deletions in logarithmic amortized time and all other heap operations in $O(1)$ amortized time. We explore the design space of this data structure. We propose a version with the following improvements over the original: (i) Each heap is represented by a single heap-ordered tree, instead of a set of trees. (ii) Each decrease-key operation does only one cut and a cascade of rank changes, instead of doing a cascade of cuts. (iii) The outcomes of all comparisons done by the algorithm are explicitly represented in the data structure, so none are wasted. We also give an example to show that without cascading cuts or rank changes, both the original data structure and the new version fail to have the desired efficiency, solving an open problem of Fredman. Finally, we illustrate the richness of the design space by proposing several alternative ways to do cascading rank changes, including a randomized strategy related to one previously proposed by Karger. We leave the analysis of these alternatives as intriguing open problems.

*Blavatnik School of Computer Science, Tel Aviv University, Israel. Research supported by The Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11). E-mail: haimk@post.tau.ac.il.

†Department of Computer Science, Princeton University, Princeton, NJ 08540 and Microsoft Research Silicon Valley, Mountain View, CA 94043.

‡Blavatnik School of Computer Science, Tel Aviv University, Israel. Research supported by BSF grant no. 2012338 and by The Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11). E-mail: zwick@tau.ac.il.

1 Introduction

A *heap* is a data structure consisting of a set of *items*, each with a key selected from a totally ordered universe. Heaps support the following operations:

`make-heap`(): Return a new, empty heap.

`find-min`(H) : Return an item of minimum key in heap H. If H is empty, return null.

`insert`(x, H): Return a heap formed from heap H by inserting item x, with predefined key. Item x must be in no heap.

`delete-min`(H): Return a heap formed from non-empty heap H by deleting the item returned by `find-min`(H).

`meld`(H_1, H_2): Return a heap containing all items in item-disjoint heaps H_1 and H_2.

`decrease-key`(x, v, H): Given that x is an item in heap H with key no less than v, return a heap formed from H by changing the key of x to v.

`delete`(x, H) : Given that x is an item in heap H, return a heap formed by deleting x from H.

The original heap H passed to `insert`, `delete-min`, `decrease-key`, and `delete`, and the heaps H_1 and H_2 passed to `meld`, are destroyed by the operations. Heaps do *not* support search by key; operations `decrease-key` and `delete` are given the location of item x in heap H as part of the input. The parameter H can be omitted from `decrease-key` and `delete`, but then to make `meld` operations efficient one needs a separate disjoint set data structure to keep track of the partition of items into heaps. For further discussion of this see [14].

Fredman and Tarjan [10] invented the *Fibonacci heap*, an implementation of heaps that supports `delete-min` and `delete` on n-item heaps in $O(\log n)$ amortized time and each of the other operations in $O(1)$ amortized time. Applications of Fibonacci heaps include a fast implementation of Dijkstra's shortest path algorithm [4, 10] and fast algorithms for undirected and directed minimum spanning trees [5, 11]. Since the invention of Fibonacci heaps, a number of other heap implementations with the same amortized time bounds have been proposed [1, 2, 3, 6, 12, 13, 15, 19, 21]. Notably, Brodal [1] invented a very complicated heap implementation that achieves the time bounds of Fibonacci heaps in the worst case. Brodal et al. [2] later simplified this data structure, but it is still significantly more complicated than any of the amortized-efficient structures. For further discussion of these and related results, see [12].

In spite of the many competitors to Fibonacci heaps, the original data structure remains one of the simplest to describe and implement. We explore the design space of this data structure. Our contributions are three. First we present a version of Fibonacci heaps in which each heap is represented by a single heap-ordered tree instead of a set of trees, each `decrease-key` is implemented using cascading rank decreases instead of cascading cuts, and the outcome of each key comparison is explicitly represented in the data structure, so that no comparisons are wasted. Second, we give an example to show that without cascading cuts or rank decreases, both the original data structure and ours fail to have the desired efficiency, solving an open problem of Fredman [8]. Finally, to illustrate the richness of the design space we propose several alternative ways to do cascading rank decreases, including a randomized method related to one previously proposed by Karger [16]. We leave the analysis of these alternatives as intriguing open problems.

The remainder of our paper consists of six sections. Section 2 describes our data structure. Section 3 analyzes it. Section 4 presents an implementation. Section 5 gives an example showing that cascading is necessary to make Fibonacci heaps efficient, answering an open question of Fredman [8]. Section 6 explores alternative ways to do cascading rank decreases. Section 7 contains final remarks.

2 Simple Fibonacci Heaps

We obtain our version of Fibonacci heaps by refining a well-known generic heap implementation. We represent a heap by a rooted tree whose nodes are the heap items. We access the tree via its root. The tree is heap-ordered: each child has key no less than that of its parent. Heap order implies that the root is an item of minimum key. To do a `find-min` on a heap, we return the root of its tree.

We do the update operations using a primitive called *linking*. Given the roots of two trees, we link them by comparing their keys and making the root of smaller key the parent of the root of the larger key, breaking a tie arbitrarily. To make a heap, we create a new, empty tree. To meld two heaps, if one is empty we return the other; if both are non-empty, we link the roots of their trees and return the resulting tree. To insert an item into a heap, we make it into a one-node tree and meld this tree with the tree representing the heap.

We do `delete-min` by repeated linking. First, we delete the root of the tree representing the heap, making each of its children into the root of a separate tree. If the root had no children, the heap is now empty. Otherwise, we repeatedly link two roots until only one tree remains, and return this tree.

To decrease the key of an item x in a heap, we begin by replacing the key of x. If x is the root of its tree, this completes the operation. Otherwise, we break the tree containing x into two trees, one containing all descendants of x (including x), the other containing the rest of the nodes in the tree. We call this a *cut* of x. We complete the operation by linking x with the root of the original tree.

To delete an arbitrary item x from its heap, we decrease its key to a value less than all other keys, and then do a `delete-min`.

The efficiency of this implementation depends on how links are done in `delete-min`. (This is the only flexibility in the implementation.) To keep the number of links small, we give each node x a non-negative integer rank, denoted by $x.rank$. We use ranks in a special kind of link called a *fair link*. A fair link of two roots can be done only if the roots have equal rank. Such a link increases the rank of the root of smaller key by 1 and makes this root the parent of the root of larger key, breaking a tie arbitrarily. In contrast, a *naïve link* ignores the ranks and merely makes the root of smaller key the parent of the root of larger key, without changing any ranks.

When linking two roots of the same rank, we have the choice of doing either a fair link or a naïve link. We do fair links only when necessary to guarantee efficiency, namely only in `delete-min` operations, since they require the extra step of changing a rank.

In addition to a rank, we give each node a state, either *unmarked* or *marked*. We represent the state of a node x by a Boolean variable $x.state$ that is true if and only if x is marked. We use Boolean constants *unmarked* and *marked* whose values are *false* and *true*, respectively. We refine the generic heap implementation as follows. Each node added by an insertion has an initial rank of 0 and is initially unmarked. Each link during an `insert` or `meld` is a naïve link. During a `delete-min`, as long as there are two roots of the same rank, we do a fair link of two such roots. Once all remaining roots have different ranks, we link them in any order by naïve links. When decreasing the key of an item x in a heap with root $h \neq x$, we unmark h and then walk up the path in the tree from the parent of x through ancestors, unmarking each marked node reached and decreasing its rank by 1 if its rank is positive, until reaching an unmarked node. We mark this node and decrease its rank by 1 if its rank is positive. Then we cut x, and link x and h via a naïve link.

To make the cascading rank change process completely transparent, we give a fragment of pseudo-code that implements it. We denote the parent of a node x by $x.parent$. The following code does the rank and

state changes, assuming $x \neq h$:

$h.state \leftarrow unmarked$
$y \leftarrow x$
`repeat`:
{ $y \leftarrow y.parent$
`if` $y.rank > 0$: $y.rank \leftarrow y.rank - 1$
$y.state \leftarrow$ `not` $y.state$ }
`until` $y.state = marked$

We call this data structure the *simple Fibonacci heap*. There are three major differences between it and the original version of Fibonacci heaps. First, the original does only fair links, not naïve links. This makes each heap a set of trees rather than just one tree, accessed via a pointer to a root of minimum key. Second, during the cascading process in a `decrease-key` operation, the original method cuts each marked node reached, as well as unmarking it and decreasing its rank. But except for cutting the node whose key decreases, these cuts are superfluous, as our analysis in the next section shows. Third, our data structure maintains the outcomes of all key comparisons in the data structure itself, via fair and naïve links. In the original, maintenance of a pointer to the root of minimum key requires comparisons among root keys whose outcomes are not explicitly maintained. Also, when decreasing the key of x, x is cut only if x is not a root and its new key is less than that of its parent. The outcome of this extra comparison is also not explicitly maintained.

3 Analysis of Simple Fibonacci Heaps

Our analysis uses the concept of *active* and *passive* children. A root that becomes a child by a fair link is active; one that becomes a child by a naïve link is passive. An active child that is marked and then unmarked becomes passive when it is unmarked; it remains passive until it becomes a root and then a child by a fair link once again. We use active and passive children in the analysis only, not in the algorithm.

Lemma 3.1 *For any node x, x has at least $x.rank$ active children.*

Proof: The only way $x.rank$ can increase is via a fair link, which adds an active child. When x loses an active child, either by a cut or because a child becomes unmarked and hence passive, $x.rank$ decreases by 1 unless it is already 0. The lemma follows by induction on time. □

Let the *size* of a node x, denoted by $x.size$, be its number of descendants, including itself. We prove that the rank of a node is at most logarithmic in its size. Recall the definition of the Fibonacci numbers: $F_0 = 0$, $F_1 = 1$, $F_k = F_{k-1} + F_{k-2}$ for $k \geq 2$. The Fibonacci numbers satisfy $F_{k+2} \geq \phi^k$, where $\phi = (1 + \sqrt{5})/2$ is the *golden ratio* [17].

Lemma 3.2 *Let n_k be the minimum size of a node of rank at least k. Then $n_k \geq \phi^k$.*

Proof: Clearly $n_0 = 1$ and $n_1 = 2$. Let x be a node of rank at least $k \geq 2$ such that $x.size = n_k$. By Lemma 3.1, x has at least k active children. Arrange these children in the order they were linked to x, most recent first. Fix $i < k$. Let y be the i^{th} active child of x in this order. We claim that just after x acquires y as a child, $x.rank \geq k - i + 1$. To prove the claim, we observe that after x acquires y as a child, $x.rank$ increases only when it acquires an active child, and of the active children acquired after y, only $i - 1$ are currently active. For every other active child acquired after y, $x.rank$ first increases and then decreases

by one, since each such child either becomes passive or is cut by a `decrease-key` operation. Hence the net increase in $x.rank$ from the time y becomes a child of x to the current time is at most $i-1$. The claim follows.

The claim implies that $y.rank \geq k-i$ when y becomes a child of x. Subsequently $y.rank$ can decrease by at most 1, or y would become passive. Hence $y.rank \geq k-i-1$, which implies $y.size \geq n_{k-i-1}$. Applying this bound to each of the first $k-1$ active children of x and adding 2 for x and its k^{th} active child, we obtain $n_k = x.size \geq n_{k-2} + n_{k-3} + \ldots + n_0 + 2$. This implies $n_k \geq n_{k-1} + n_{k-2}$ for $k \geq 2$. It follows by induction on k that $n_k \geq F_{k+2} \geq \phi^k$. □

Corollary 3.3 *A node of size n has rank at most $\log_\phi n$.*

To analyze the amortized efficiency of the heap operations, we use the standard potential-based framework [22]. We assign to each configuration of the data structure a real-valued *potential*. We define the amortized time of an operation to be its actual time (possibly scaled by a constant factor) plus the net increase in potential caused by the operation. With this definition, the total time of a sequence of operations is equal to the sum of the amortized times of the operations plus the initial potential minus the final potential. If the initial potential is zero and the final potential is non-negative, then the total actual time of the operations is at most the sum of their amortized times (times the scaling constant if any).

In our analysis we estimate the time of an operation other than a `delete-min`, `decrease-key`, or `delete` to be 1; that of a `decrease-key` to be $1 + \#iterations$, where $\#iterations$ is the number of iterations of the rank-decreasing loop; and that of a `delete-min` to be $1 + \log_\phi n + \#links$, where n is the number of items in the heap and $\#links$ is the number of links, both fair and naïve, done during the operation. We treat a `delete` as a `decrease-key` followed by a `delete-min`. In the next section we present an implementation of the data structure that supports each operation except `delete-min`, `decrease-key`, and `delete` in $O(1)$ actual time, each `decrease-key` in $O(1 + \#iterations)$ actual time, and each `delete-min` in $O(1 + \log_\phi n + \#links)$ actual time, justifying our estimates.

Let the *degree* of a node x, denoted by $x.degree$, be the number of children of x. We use degrees in the analysis only; they are not part of the algorithm. We define the potential of a node x to be $x.degree - x.rank$, plus 1 if x is a root, plus 2 if x is marked. For any node x, $x.degree \geq x.rank$ by Lemma 3.1, so every node has non-negative potential, and every root has potential at least 1. We define the potential of a set of heaps to be the sum of the potentials of their nodes. This potential is 0 initially (there are no nodes) and is always non-negative. Thus in any sequence of heap operations the sum of the amortized times is an upper bound on the sum of their estimated times.

Lemma 3.4 *The amortized time of a* `make-heap`, `find-min`, *or* `meld` *is* 1, *that of an* `insert` *is* 2, *that of a* `decrease-key` *is at most* 5, *and that of a* `delete-min` *on a heap of n items is at most $3\log_\phi n$.*

Proof: A naïve link does not change the potential: the potential of the root that becomes a child decreases by 1 but that of the root that remains a root increases by 1 (its degree increases and its rank does not change). A fair link decreases the potential by 1: in this case the potential of the root that remains a root does not change, since both its degree and its rank increase by 1. A `make-heap`, `find-min`, or `meld` does not change the potential, so the amortized time of such an operation equals its estimated time of 1. An insert creates a new root, of potential 1, making its amortized time 2 whether or not it does a naïve link.

A `decrease-key` of a root does not change the potential and hence has an amortized time of 1. Consider a `decrease-key` of a non-root x. If the root is marked, unmarking it decreases the potential. Each iteration of the rank-decreasing loop except the last unmarks a node y and may decrease its rank by 1, decreasing the potential by at least $2 - 1 = 1$. The last iteration marks a node and may decrease its rank by one, increasing the potential by at most 3. Cutting x from its parent does not change the potential:

the potential of x increases by 1 since it becomes a root, but that of its parent decreases by 1, since its degree decreases by 1 and its rank does not change. (Its rank may have decreased by 1 earlier, but we have already accounted for this.) Combining these observations, we find that the `decrease-key` increases the potential by at most $3 - (\#iterations - 1)$, making the amortized time of the operation at most $(1 + \#iterations) + (4 - \#iterations) = 5$.

Finally, consider a `delete-min` on a heap of $n > 0$ items. Let h be the root of the tree representing the heap. Deleting h increases the potential by at most $h.rank - 1$: deletion of h reduces the potential by at least $1 + h.degree - h.rank$, but each of the $h.degree$ new roots increases in potential by 1. By Corollary 3.3, $h.rank \leq \log_\phi n$, so deleting h increases the potential by at most $\log_\phi n - 1$. Each fair link reduces the potential by 1. Also by Corollary 3.3, at most $\log_\phi n$ of the links are naïve, so at least $\#links - \log_\phi n$ are fair. We conclude that the entire `delete-min` increases the potential by at most $\log_\phi n - 1 + \log_\phi n - \#links$, making the amortized time of the operation at most $(1 + \log_\phi n + \#links) + (2\log_\phi n - 1 - \#links) = 3\log_\phi n$. □

Theorem 3.5 *The total estimated time of a sequence of heap operations starting with no heaps is $O(1)$ per operation other than* `delete-min` *and* `delete`*, and $O(\log n)$ per* `delete-min` *and* `delete`*, where n is the number of items currently in the heap.*

Proof: The theorem is immediate from Lemma 3.4. □

Remark: The potential used in the proof of Lemma 3.4 is like the original [10] except that potential associated with children is moved to their parents. This is necessary because Lemma 3.1 is an inequality, not an equality. Alternatively, we can define a child y of a node x to be *hyperactive* if y is one of the $x.rank$ most recently linked active children of x. Then the number of hyperactive children of a node equals its rank; and, if we give each node 1 unit of potential if it is not a hyperactive child, plus 2 if it is marked, then Lemma 3.4 holds. (The potential of a tree is the same under both definitions.)

4 Implementation

Our implementation of simple Fibonacci heaps mimics the original implementation [10]. We make each set of children a doubly-linked list: if x is a node, $x.after$ and $x.before$ are its next sibling and its previous sibling, respectively. We need double linking to do cuts in $O(1)$ time. We also store with each node x pointers to its parent, $x.parent$, and to its first child on its list of children, $x.child$. To find fair links to do in `delete-min`, we use a global array A of nodes indexed by rank. We assume that the array is initially empty. After deleting the root, we place each new root into the array location corresponding to its rank. If this location is already occupied, we do a fair link. Once all remaining roots have been successfully stored in the array, we scan the array, linking pairs of roots by naïve links, until only one root remains, and leave the array empty again. This implementation supports all the operations in the time bounds claimed in Section 3. Since the rank of a node increases or decreases by only 1 at a time, we can replace the global array by a global doubly linked list to obtain an implementation on a pointer machine. See [10]. Figure 1 gives pseudo-code implementing the heap operations, as well as `make-item`$(info, v)$, which creates a heap item having key v and associated information *info*. Figure 2 gives pseudo-code implementing the auxiliary functions `link`, `cut`, `add-child`, and `decrease-ranks` used in Figure 1.

There are several generic changes that can be made to our implementation, or to an implementation of any similar data structure for heaps. One change is to reduce the number of pointers per node from 4 to 3 by replacing the *before* and *after* pointers by one pointer whose meaning depends on the type of node: a first child points to the next sibling of its parent; a child other than the first points to its previous sibling. (See [9].) This trades time for space. More complicated representations allow further space reductions.

```
make-item(info, v):
{   x ← new node
    x.info ← info
    x.key ← v
    x.rank ← 0
    x.state ← unmarked
    x.child ← null
    return x  }

make-heap():
    return null

insert(x, h):
    return meld(x, h)

meld(g, h):
{   if g = null: return h
    if h = null: return g
    return link(g, h)  }

delete(x, h):
{   decrease-key(x, −∞, h)
    return delete-min(h)  }
```

```
delete-min(h):
{   x ← h.child
    max-rank ← 0
    while x ≠ null:
    {   y ← x
        x ← x.after
        while A[y.rank] ≠ null:
        {   y ← link(y, A[y.rank])
            A[y.rank] ← null
            y.rank ← y.rank + 1  }
        A[y.rank] ← y
        if y.rank > max-rank : max-rank ← y.rank  }
    for i ← 0 to max-rank:
    {   if A[i] ≠ null:
        {   if x = null: x ← A[i]
            else: x ← link(x, A[i])
            A[i] ← null  }  }
    return x  }

decrease-key(x, v, h):
{   x.key ← v
    if x = h: return h
    h.state ← unmarked
    decrease-ranks(x)
    cut(x)
    return link(x, h)  }
```

Figure 1: Implementation of the heap operations

```
link(x, y):
{   if x.key > y.key:
    {   add-child(x, y)
        return y  }
    else:
    {   add-child(y, x)
        return x  } }

add-child(x, y):
{   x.parent ← y
    z ← y.child
    x.before ← null
    x.after ← z
    if z ≠ null: z.before ← x
    y.child ← x  }
```

```
cut(x):
{   y ← x.parent
    if y.child = x: y.child ← x.after
    if x.before ≠ null: x.before.after ← x.after
    if x.after ≠ null: x.after.before ← x.before  }

decrease-ranks(y):
{   repeat:
    {   y ← y.parent
        if y.rank > 0: y.rank ← y.rank − 1
        y.state ← not y.state  }
    until y.state = marked  }
```

Figure 2: Auxiliary functions used in the implementation of the heap operations

A second change is to modify the specification of the heap operations so that they update old heaps instead of returning new ones. One can do this by introducing separate heap objects. In a minimal implementation

```
Make-heap():
{   H ← new heap
    H.root ← null
    return H  }

Meld(H1, H2):
    H1.root ← meld(H1.root, H2.root)

Insert(x, H):
    H.root ← meld(x, H.root)

Find-min(H):
    return H.root

Delete-min(H):
    H.root ← delete-min(H.root)

Decrease-key(H, x, k):
    H.root ← decrease-key(x, k, H.root)

Delete(H, x) :
    H.root ← delete(x, H.root)
```

Figure 3: An implementation that modifies existing heaps.

of this idea, a heap object H has a single field $H.root$ that is the root of the tree representing the heap, or null if the heap is empty. A less minimal implementation would store other useful information, such as the heap size, in each heap object. Figure 3 gives an implementation using heap objects. The operation $Insert(x, H)$ replaces $h \leftarrow insert(x, h)$, and similarly for the other operations.

A third change is to modify the implementation to do links lazily, only when needed to answer `find-min` queries. With this change, a heap consists of a set of heap-ordered trees, represented by a list of roots. A `meld` catenates the lists representing the two input heaps, without doing any links. A `decrease-key` of a non-root node x adds to the list of roots instead of doing a link. A `delete-min` does a `find-min`, which converts the list of roots into a single root, and then deletes the root, leaving its list of children as the new list of roots. A `find-min` links roots as in the implementation of `delete-min` in Section 2, but without first deleting a node; once all roots are linked, it returns the one remaining root. The $O(1)$ time bound for `find-min` becomes amortized instead of worst-case. As mentioned in Section 3, the original version of Fibonacci heaps also represents the heap by a list of roots, but unlike the lazy implementation just described it maintains a pointer to a root of minimum key (making the time for `find-min` $O(1)$ worst-case). Also, it does only fair links, and it does all its links in `delete-min`, not `find-min`.

These three changes are independent of each other, so they can be combined arbitrarily.

We conclude this section by mentioning three heuristics that can be added to `decrease-key` to reduce the number of iterations of the rank-decreasing loop in certain cases. These heuristics are independent and hence can be combined arbitrarily. Whether they provide improved performance is a question for experiments; none improves the $O(1)$ amortized time bound for `decrease-key`, and they all complicate the implementation at least slightly, but they do preserve the analysis of Section 3. Furthermore, they provide ideas for actually simplifying the rank-decreasing process that we explore in Section 6.

The *heap-order heuristic* cuts x in `decrease-key`(x, v, h) only if x is not a root *and* its new key, v, is less than that of its parent, violating heap order. As noted in Section 2, the original version of Fibonacci heaps uses this heuristic. A possible disadvantage of this heuristic is that the outcomes of such comparisons are not encoded in the data structure.

We can combine the two tests of whether x is a child and whether x violates heap order into a single test by defining the parent of a root to be itself. Then $x.key < x.parent.key$ if and only if x is a child that violates heap order. Maintaining parents of roots requires initializing them in `make-item` and updating them in `cut`. In contrast, the implementation in Figures 1 and 2 does not need to maintain parents of roots, since it never uses them.

The *increasing-rank heuristic* stops the rank-decreasing loop when it reaches a node whose old rank is no less than that of its parent. This heuristic reduces the worst-case time of `decrease-key` to $O(\log n)$, since nodes visited in the rank-decreasing loop strictly increase in old rank. To implement the heuristic, replace the implementations of `decrease-key` and `decrease-ranks` by the implementations in Figure 4.

```
decrease-key(x, v, h):
{  x.key ← v
   if x = h: return h
   if x.rank < x.parent.rank:
   {  h.state ← unmarked
      decrease-ranks(x)  }
   cut(x)
   return link(x, h)  }
```

```
decrease-ranks(y):
{  repeat:
   {  y ← y.parent
      y.state ← not y.state
      k ← y.rank
      y.rank ← k − 1  }
   until y.state = marked or k ≥ y.parent.rank  }
```

Figure 4: Implementation of `decrease-key` using the increasing-rank heuristic.

The *passive child heuristic* stops the rank-decreasing loop when it reaches a passive child (as defined in Section 3). To implement this heuristic, give each node one of three possible states: *passive*, *unmarked* (and active), or *marked* (and active). Make roots implicitly passive. Modify `link` to make the new child passive, modify `delete-min` to make each child added by a fair link unmarked, and replace the implementations of `decrease-key` and `decrease-ranks` by the implementations in Figure 5.

```
decrease-key(x, v, h):
{  x.key ← v
   if x = h: return h
   h.state ← passive
   if x.state ≠ passive: decrease-ranks(x)
   cut(x)
   return link(x, h)  }
```

```
decrease-ranks(y):
{  while y.state = marked:
   {  y.state ← passive
      y.rank ← y.rank − 1
      y ← y.parent  }
   y.rank ← y.rank − 1
   if y.state = unmarked: y.state ← marked  }
```

Figure 5: Implementation of `decrease-key` using the passive child heuristic.

The passive child heuristic provides a bridge from the simple Fibonacci heaps of Section 2 to the original version [10]. Indeed, we originally included this heuristic in our data structure, and then omitted it to obtain the version in Section 2. Simple Fibonacci heaps with the heap-order and passive-child heuristics correspond exactly to the original version of Fibonacci heaps. To obtain the data structure of the latter from that of the former, cut each passive child, producing a forest. To map operations of the former to those of the latter, omit all the naïve links, and for each node y unmarked by `decrease-ranks`, cut y. With the potential function redefined as in the remark at the end of Section 3, the analysis of the former corresponds exactly to that of the latter.

5 Cascading is Required

We have replaced the cascading cuts of Fibonacci heaps with cascading rank decreases. Fredman [8] has asked whether one can avoid cascading altogether. Specifically, consider the following simplification of our data structure: keep ranks but eliminate node marking. To decrease the key of a non-root x, merely cut x from its parent, decrease the rank of the parent by 1 if its rank is positive, and link x and the original root by a naïve link. We call this data structure a *non-cascading heap*. We show how to build a non-cascading heap of n nodes on which one `insert` followed by one `delete-min` reproduces the original structure, with the `delete-min` taking $\Omega(n^{1/2})$ time. Such pairs of `insert` and `delete-min` operations can be repeated indefinitely. The number of operations needed to build the heap is $O(n^{3/2})$. It follows that a sequence of m operations starting with no heaps can take $\Omega(m^{4/3})$ time.

Define S_k for a non-negative integer k to be a tree of $k+1$ nodes consisting of a root of rank k with k

children. Define $T_k(i)$ for a positive integer k and a non-negative integer $i \leq k$ to be a tree whose root has k children that are the roots of $S_0, S_1, \ldots, S_k$, with S_i missing. Thus $T_k(i)$ for $0 < i < k$ consists of a root whose k children are the roots of $S_0, \ldots, S_{i-1}, S_{i+1}, \ldots, S_k$. See Figure 6. As a special case, $T_k(k)$ consists of a root whose k children are the roots of $S_0, S_1, ..., S_{k-1}$. Tree $T_k(k)$ has size $1 + k(k-1)/2$. Given a heap whose tree is $T_k(k)$, one `insert` of an item with key larger than that of the root but smaller than that of all other nodes in the tree, followed by one `delete-min`, will produce a new copy of $T_k(k)$, the `delete-min` having done k fair links and spent $\Omega(k)$ time. Thus to get a bad example all we need to do is build a $T_k(k)$ for large k.

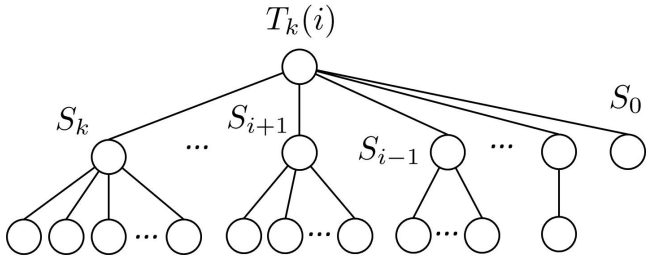

Figure 6: The tree $T_k(i)$, for $1 < i < k$.

We can convert $T_k(0)$ into $T_{k+1}(k+1)$ by inserting a new item whose key is greater than that of the root. We can convert $T_k(i)$ with $i > 0$ into $T_k(i-1)$ by doing two insertions, one `delete-min`, and $i-1$ `decrease-key` operations, as follows. We insert two items whose keys are larger than that of the root. Then we do a `delete-min`. After deletion of the original root, there are three roots of rank 0: the two newly inserted items and the root of the original S_0. We link one of the newly inserted nodes with the root of the original S_0 by a fair link and then successively link the resulting tree with S_1, S_2, ..., S_{i-1} by fair links. We choose the keys so that among the roots linked, the root of S_{i-1} has smallest key and one of the newly inserted nodes has second-smallest key. Then the tree formed by the links is S_i but with one of its leaves having children that are the roots of S_0, S_1,..., S_{i-2}. We then combine the remaining trees, which include one singleton, by naïve links that make the singleton the final root and all the other roots its children. Finally we do $i-1$ `decrease-key` operations on the roots of S_0, S_1,..., S_{i-2}, making their roots into children of the final root. This produces $T_k(i-1)$. See Figure 7.

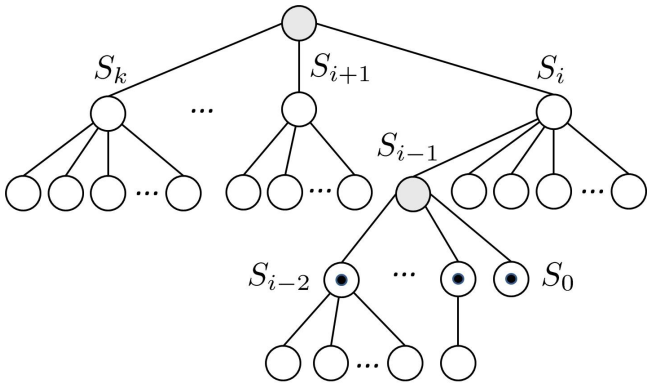

Figure 7: The tree obtained when transforming $T_k(i)$ to $T_k(i-1)$, after two insertions and a `delete-min`. The last step, which is not shown, is to do $i-1$ `decrease-key` operations on the roots of S_0, S_1,..., S_{i-2} marked with a dot.

By induction, we can convert $T_k(k)$ into $T_{k+1}(k+1)$ in $O(k^2)$ heap operations, and we can build $T_k(k)$ from scratch in $O(k^3)$ heap operations.

Theorem 5.1 *For any n, there is a non-cascading heap of n nodes on which an* `insert` *followed by a* `delete-min` *can take $\Omega(n^{1/2})$ time and reproduce the original structure. Such a heap can be built from scratch in $O(n^{3/2})$ operations. For any m there is a sequence of m operations starting from no heaps that takes $\Omega(m^{4/3})$ time.*

Proof: The first two parts of the theorem are immediate from the construction above. To prove the third part, assume $m \geq 6$. Let k be maximum such that T_k can be built in at most $m/3$ operations. Then $k = \Omega(m^{1/3})$ by the construction above. Build T_k in at most $m/3$ operations and then alternate `insert` and `delete-min` operations, with each pair recreating T_k and taking $\Omega(k)$ time. □

Essentially the same example shows that if cascading cutting is eliminated from the original version of Fibonacci heaps, the resulting data structure can take fractional-polynomial time per operation. This shows that the answer to Fredman's question is "No".

6 Simpler Cascading?

Given that cascading rank decreases are necessary, it is natural to ask whether there is a simpler way to decide when to do them. We offer four possible methods, three deterministic and one randomized. We cannot yet fully analyze any of these methods, and we leave doing so as intriguing open problems.

Two of our three deterministic methods simplify two of the heuristics in Section 4. *Eager marking* simplifies the passive child heuristic by making marked nodes and active nodes identical. Roots are implicitly unmarked. A node that becomes a child by a fair link becomes marked. The rank-decreasing loop stops when it reaches an unmarked node, which it leaves unmarked. To implement this method, modify `link` to make the new child unmarked, modify `delete-min` to make each child added by a fair link marked, and replace the implementation of `decrease-ranks` by the implementation in Figure 8, given in the appendix.

Our other two deterministic methods eliminate marking entirely. The *naïve increasing-rank method* stops the rank-decreasing loop when it reaches a node whose old rank is no less than that of its parent. The *zero-rank method* stops the rank-decreasing loop when it reaches a node whose rank is 0. To implement these methods, eliminate node states and replace the implementations of `decrease-key` and `decrease-ranks` by those in Figures 9 and 10, respectively, given in the appendix.

It is straightforward to prove that with each of these methods, a node of size n has rank at most $\lg n$. It follows that the worst-case time of `decrease-key` with the naive increasing-rank method is $O(\log n)$, and that the bounds of Theorem 3.5 hold for this method except that the amortized time per `decrease-key` is $O(\log n)$. For the other two methods we cannot prove even this weaker result. We conjecture that Theorem 3.5 holds for none of these methods, but we have no counterexamples.

Another way to eliminate marking is to use *randomized cascading*, which at each iteration of the rank-decreasing loop stops with probability 1/2. To implement this method, eliminate node stated and replace the implementations of `decrease-key` and `decrease-ranks` by those in Figure 11, given in the appendix.

Randomized cascading is a variant of a method proposed by Karger [16], *randomized cutting*, which applies to the original version of Fibonacci heaps [10]. It does not use marking; instead, at each iteration of the cascading-cut loop it stops with probability 1/2. With either randomized cascading or randomized cutting, `decrease-key` takes $O(1)$ expected time. In recent work independent of ours, Li and Peebles [18] have given a tight analysis of randomized cutting. They showed that the expected amortized time of `delete-min` is $\Theta(\min\{n^{1/2}, (\log n \log s)/\log\log s\})$, where s is the total number of `decrease-key` operations. Thus randomized cutting does not provide the efficiency of Fibonacci heaps. Their lower bound example is similar to our example in Section 5. Their upper bound analysis applies in a straightforward way to randomized cascading, but their lower bound does not (as far as we can tell). Thus it remains open whether randomized cascading achieves the efficiency of Fibonacci heaps (in expectation).

7 Remarks

We have presented a one-tree version of Fibonacci heaps in which each `decrease-key` requires only a single cut. Ours is not the only heap structure with these properties. For example, rank-pairing heaps [12] were specifically designed to need only one cut per `decrease-key`, and there is a one-tree version of these heaps [12]. But among all the known data structures that have the same amortized time bounds as Fibonacci heaps, we think ours is the simplest. It has the additional nice property that the outcomes of all key comparisons are represented in the data structure (via links). The pairing heap [9], a simpler self-adjusting structure, is faster in practice, but it does not support `decrease-key` in $O(1)$ amortized time [7]. Indeed, obtaining a tight amortized time bound for `decrease-key` in pairing heaps is an interesting open problem [7, 20].

We have shown that without cascading rank changes, our data structure does not have the efficiency of Fibonacci heaps, solving an open problem of Fredman [8]. We have proposed four simplified cascading methods. We leave the analysis of these methods as intriguing open problems.

References

[1] G.S. Brodal. Worst-case efficient priority queues. In *Proceedings of the 7th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 52–58, 1996.

[2] G.S. Brodal, G. Lagogiannis, and R.E. Tarjan. Strict Fibonacci heaps. In *Proceedings of the 44th ACM Symposium on Theory of Computing (STOC)*, pages 1177–1184, 2012.

[3] T.M. Chan. Quake heaps: A simple alternative to Fibonacci heaps. In *Space-Efficient Data Structures, Streams, and Algorithms*, pages 27–32, 2013.

[4] E.W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.

[5] J. Edmonds. Optimum branchings. *J. Res. Nat. Bur. Standards*, 71B:233–240, 1967.

[6] A. Elmasry. The violation heap: a relaxed Fibonacci-like heap. *Discrete Math., Alg. and Appl.*, 2(4):493–504, 2010.

[7] M.L Fredman. On the efficiency of pairing heaps and related data structures. *Journal of the ACM*, 46(4):473–501, 1999.

[8] M.L. Fredman. Binomial, Fibonacci, and pairing heaps. In D.P. Mehta and S. Sahni, editors, *Handbook of data structures and applications*, pages 7-1–7-20. Chapman & Hall/CRC, Boca Raton, 2005.

[9] M.L. Fredman, R. Sedgewick, D.D. Sleator, and R.E. Tarjan. The pairing heap: a new form of self-adjusting heap. *Algorithmica*, 1(1):111–129, 1986.

[10] M.L. Fredman and R.E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the ACM*, 34(3):596–615, 1987.

[11] H.N. Gabow, Z. Galil, T.H. Spencer, and R.E. Tarjan. Efficient algorithms for finding minimum spanning trees in undirected and directed graphs. *Combinatorica*, 6:109–122, 1986.

[12] B. Haeupler, S. Sen, and R.E. Tarjan. Rank-pairing heaps. *SIAM Journal on Computing*, 40(6):1463–1485, 2011.

[13] P. Høyer. A general technique for implementation of efficient priority queues. In *Proceedings of the 3rd Israeli Symposium on the Theory of Computing and Systems (ISTCS)*, pages 57–66, 1995.

[14] H. Kaplan, N. Shafrir, and R.E. Tarjan. Meldable heaps and boolean union-find. In *Proceedings of the 34th ACM Symposium on Theory of Computing (STOC)*, pages 573–582, 2002.

[15] H. Kaplan and R.E. Tarjan. Thin heaps, thick heaps. *ACM Transactions on Algorithms*, 4(1):1–14, 2008.

[16] D. Karger. Private communication to D. Roy, 2006.

[17] D.E. Knuth. *Sorting and searching*, volume 3 of *The art of computer programming*. Addison-Wesley, second edition, 1998.

[18] J. Li and J. Peebles. Replacing mark bits with randomness in Fibonacci heaps. Technical report, 2014. arXiv:1407.2569v2.

[19] G.L. Peterson. A balanced tree scheme for meldable heaps with updates. Technical Report GIT-ICS-87-23, School of Informatics and Computer Science, Georgia Institute of Technology, Atlanta, GA, 1987.

[20] S. Pettie. Towards a final analysis of pairing heaps. In *Proceedings of 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 174–183, 2005.

[21] T. Takaoka. Theory of 2-3 heaps. *Discrete Applied Mathematics*, 126(1):115–128, 2003.

[22] R.E. Tarjan. Amortized computational complexity. *SIAM Journal on Algebraic and Discrete Methods*, 6:306–318, 1985.

```
decrease-ranks(y):
    while y.state = marked:
    {   y.state ← unmarked
        y ← y.parent
        y.rank ← y.rank − 1  }
```

Figure 8: Implementation of `decrease-ranks` using eager marking.

```
decrease-key(x, v, h):
{   x.key ← v
    if x = h: return h
    if x.rank < x.parent.rank: decrease-ranks(x, h)
    cut(x)
    return link(x, h)  }
```

```
decrease-ranks(y, h):
{   repeat:
        {   y ← y.parent
            k ← y.rank
            y.rank ← k − 1  }
        until y = h or k ≥ y.parent.rank  }
```

Figure 9: Implementation of `decrease-key` using the naïve decreasing-rank method.

```
decrease-key(x, v, h):
{   x.key ← v
    if x = h: return h
    if x.parent.rank > 0: decrease-ranks(x, h)
    cut(x)
    return link(x, h)  }
```

```
decrease-ranks(y, h):
{   repeat
    {   y ← y.parent
        y.rank ← y.rank − 1  }
    until y = h or y.parent.rank = 0  }
```

Figure 10: Implementation of `decrease-key` using the zero-rank method.

```
decrease-key(x, v, h):
{   x.key ← v
    if x = h: return h
    decrease-ranks(x)
    cut(x)
    return link(x, h)  }
```

```
decrease-ranks(y):
{   repeat:
    {   if y.rank > 0: y.rank ← y.rank − 1
        y ← y.parent  }
    until y = h or heads()  }
```

Figure 11: Implementation of `decrease-key` using randomization. Each call of function `heads`() returns a Boolean value that is true with probability 1/2 and false with probability 1/2, independent of the values returned by previous calls.